EgoGen: An Egocentric Synthetic Data Generator

Gen Li[1] Kaifeng Zhao[1] Siwei Zhang[1] Xiaozhong Lyu[1]
Mihai Dusmanu[2] Yan Zhang[1] Marc Pollefeys[1,2] Siyu Tang[1]
[1]ETH Zürich [2]Microsoft
https://ego-gen.github.io/

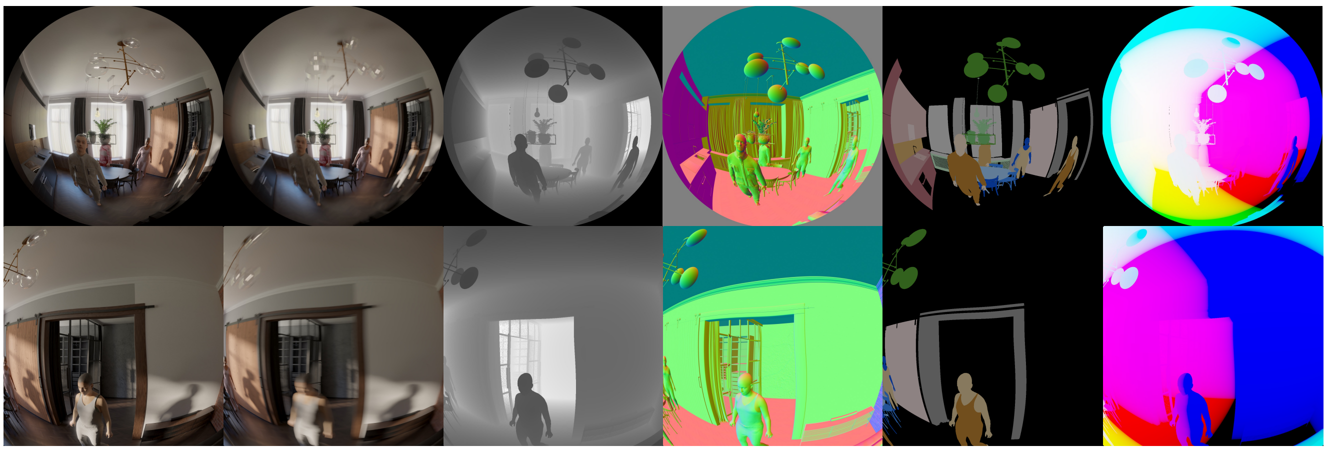

Figure 1. EgoGen: a scalable synthetic data generation system for egocentric perception tasks, with rich multi-modal data and accurate annotations. We simulate camera rigs for head-mounted devices (HMDs) and render from the perspective of the camera wearer with various sensors. Top to bottom: middle and right camera sensors in the rig. Left to right: photo-realistic RGB image, RGB with simulated motion blur, depth map, surface normal, segmentation mask, and world position for fisheye cameras widely used in HMDs.

Abstract

Understanding the world in first-person view is fundamental in Augmented Reality (AR). This immersive perspective brings dramatic visual changes and unique challenges compared to third-person views. Synthetic data has empowered third-person-view vision models, but its application to embodied egocentric perception tasks remains largely unexplored. A critical challenge lies in simulating natural human movements and behaviors that effectively steer the embodied cameras to capture a faithful egocentric representation of the 3D world. To address this challenge, we introduce EgoGen, a new synthetic data generator that can produce accurate and rich ground-truth training data for egocentric perception tasks. At the heart of EgoGen is a novel human motion synthesis model that directly leverages egocentric visual inputs of a virtual human to sense the 3D environment. Combined with collision-avoiding motion primitives and a two-stage reinforcement learning approach, our motion synthesis model offers a closed-loop solution where the embodied perception and movement of the virtual human are seamlessly coupled. Compared to previous works, our model eliminates the need for a pre-defined global path, and is directly applicable to dynamic environments. Combined with our easy-to-use and scalable data generation pipeline, we demonstrate EgoGen's efficacy in three tasks: mapping and localization for head-mounted cameras, egocentric camera tracking, and human mesh recovery from egocentric views. EgoGen will be fully open-sourced, offering a practical solution for creating realistic egocentric training data and aiming to serve as a useful tool for egocentric computer vision research.

1. Introduction

The analysis of visual input from front-facing egocentric cameras is crucial for applications that benefit from a first-person perspective, mirroring the natural human experience [23, 43, 115]. AR devices, for instance, can utilize this viewpoint to amplify user immersion. Such cameras can also cater to individual preferences, providing custom visual assistance for those with impaired vision [19, 120].

Despite its potential, egocentric perception faces challenges, primarily due to the scarcity of labeled data. Although datasets like Ego4D [23], ADT [61], Epic-Kitchen [13] and HoloAssist [102] exist, creating such datasets with rich and accurate annotations is costly and raises privacy concerns [115]. Alternatively, using graphics techniques to render synthetic multi-modal visual data has proven to be cost-effective and successful in training deep learning models, such as 3D human body estimation [8] and facial landmark detection [106].

Creating egocentric synthetic data is challenging because egocentric cameras capture the complex interplay of body movements and the environment from the camera wearer's viewpoint. Modeling the intricate details and variations in human behavior presents a significant challenge.

To tackle this problem, we introduce *EgoGen*, an egocentric synthetic data generation approach that simulates data from embodied sensors, i.e., front-facing cameras in head-mounted devices (HMD). While the ultimate goal is to simulate human behaviors that are indistinguishable from reality, in this work, we focus on creating virtual humans (i.e., camera wearers) that can explore and avoid obstacles in the 3D world that is not only complex and dynamic but could potentially include other *moving* virtual humans.

Specifically, we propose a novel generative human motion model. Our key insight is that body movement and embodied perception should be seamlessly coupled. As William Gibson aptly stated, *"We see in order to move; we move in order to see."*, our egocentric perception is crucial for identifying obstacles, navigating in an environment, and planning actions. Our body movements are not solely a response to visual stimuli; they also change our egocentric perception. Therefore, the key idea of our motion model is to enable virtual humans to *see* their environment with *egocentric* visual inputs and respond accordingly by learning a policy to control a set of collision-avoiding motion primitives (CAMPs) that are composable for synthesizing long-term, diverse human motions. Due to the unbounded and high-dimensional latent action space of our generative motion primitive model, direct policy training through rendered egocentric images is often unstable [122]. Therefore, we propose a two-stage reinforcement learning scheme using an efficient *egocentric* visual proxy to couple egocentric visual cues and body movements seamlessly. In addition, we use an "attention" reward to incentivize egocentric perception behaviors, i.e., looking in the desired direction.

Empirical results showcase the benefits of our egocentric perception-driven motion framework, which does not require a pre-calculated walking path in 3D scenes as in [29, 53, 119]. Instead, it empowers virtual humans to perceive the environment from their own viewpoint, enabling them to navigate, circumvent obstacles, and plan movements to reach the destination. Moreover, our model generalizes well to dynamic environments, even with training limited to static settings. By training virtual humans independently using CAMPs, our method synthesizes emergent multi-human behaviors without relying on multi-agent reinforcement learning algorithms. Egocentric visual cues are essential to build exploratory and generalizable motion models that unify navigational planning and movement control in complex and dynamic environments.

Building upon CAMPs, we further create a scalable data generation pipeline for *EgoGen* that outfits virtual humans with clothing, automates cloth animation, and integrates 3D assets from various sources. We validate *EgoGen*'s efficacy across three egocentric perception tasks. The high-quality synthetic data with precise ground truth annotations consistently improve the performance of state-of-the-art methods. In summary, the contributions of this work are:

1. We introduce *EgoGen*, a generative and scalable synthetic data generation approach specifically tailored for egocentric perception tasks.
2. We introduce novel motion primitives based on egocentric visual cues, enabling diverse and realistic human motion synthesis in 3D scenes. These primitives empower virtual humans to handle complex scenarios, such as dynamic environments and crowd motion without relying on multi-agent reinforcement learning.
3. *EgoGen* enables us to augment existing real-world egocentric datasets with synthetic images. Quantitative results demonstrate enhanced performance in state-of-the-art algorithms on mapping and localization for head-mounted cameras, egocentric camera tracking, and human mesh recovery from egocentric views.

2. Related Work

Human-Related Simulators and Synthetic Data. Previous works primarily focus on simulating robots [54, 69, 78, 92, 95] and autonomous cars [7, 16, 20, 74]. While some incorporated animated digital humans, like in [5, 20, 67], these efforts relied on pre-recorded motion sequences. Rendering images of people to train perception models has been widely studied such as [5, 17, 18, 43, 68, 82, 86]. In particular, [97] offers large-scale synthetic data for egocentric camera wearer pose estimation but relies on mocap data, lacking realistic and spontaneous interactions with the digital world. In contrast, *EgoGen* closes the gap in egocentric synthetic data generation for head-mounted devices. Please refer to Supp. Mat. for detailed comparisons.

Human Motion Synthesis. Generating high-fidelity human motions and interactions with 3D scenes is widely studied in graphics [10, 31, 32, 42, 89, 90]. While they can generate high-quality motion, it's usually deterministic. Synthesizing physically plausible human motions has been extensively studied [12, 30, 64, 73, 96, 105]. However, they struggle with generalization to different body shapes. For example, [73] explicitly created 2048 humanoids to improve body shape generalizability. Time series models [65, 94, 114] synthesize the stochastic motions of diverse people well. However, in [94, 114], their generated motion sequences have limited lengths and human-scene interactions are not explicitly considered. Autoregressive methods [45, 72, 117, 119] can produce perpetual motions. In particular, [117] can generalize to diverse body shapes and synthesize long-term human motions.

Our egocentric perception-driven motion synthesis model is closely related to [117, 119], but distinguishes itself w.r.t.: (1) Enabling virtual humans to explore using egocentric visual cues, without predefined paths. (2) Synthesizing egocentric perception behaviors beyond locomotion, e.g., looking in certain directions. (3) Handling dynamic environments and multi-agent behavior without re-training.

Mapping and Localization for AR. Localization and mapping from images is a long-standing problem known as: Photogrammetry [2, 26]; Structure-from-Motion (SfM) [22, 66, 79, 87, 104]; Simultaneous Localization and Mapping (SLAM) [14, 37, 56, 60]. Researchers have worked to make SLAM amenable for edge hand-held or head-mounted devices [6, 21, 37]. Cloud-based services like Google's Visual Positioning System [71], Niantic's Lightship [59], and Microsoft's Azure Spatial Anchors [34] have made visual localization and mapping more accessible. Benchmarking efforts have arisen for small-scale AR scenarios [36, 85], touristic landmarks [33, 80], and large-scale AR-device based localization [76, 77] to evaluate these systems.

Egocentric Human Pose Estimation. Estimating 3D bodies from RGB images is widely studied from third-person views [9, 35, 38–41, 44, 63], and egocentric views [25, 46, 47, 57, 83, 97, 98, 100, 112, 113, 116], mostly requiring expensive real-world data paired with ground truth annotations. Besides RGB images, depth images offer explicit 3D information, mitigating scale and shape ambiguity, with the potential to enable broader AR/VR applications. However, depth-based methods, especially for the egocentric view, are underexplored. Most existing works [28, 49, 55, 70, 84, 101, 107, 111] predict 3D body skeletons without expressive body meshes, struggling with challenges like severe body truncations and scene occlusions typical in egocentric views. Such limited attention mainly stems from the scarcity of data, as obtaining high-quality human mesh annotations for real-world depth images is labor-intensive.

3. Ego-Sensing Driven Motion Synthesis

To close the loop for the interdependence between egocentric synthetic image data and human motion synthesis, we use deep reinforcement learning (RL), integrating egocentric vision cues to synthesize human motions as described in Sec. 3.1 and 3.2. Subsequently, we extend learned policies to generate emergent multi-agent behaviors, as in Sec. 3.3.

3.1. Ego-Sensing Driven Motion Primitives

Generating realistic egocentric data requires diverse and lifelike human motion synthesis. In this work, we consider arguably the most common everyday behaviors: navigating towards goals with egocentric perception while avoiding collisions with obstacles and people in dynamic 3D scenes.

Overview. Following recent literature [45, 110, 117, 119], we employ deep RL to train control policies on learned latent spaces that characterize natural human motions. However, unlike these previous works that only consider simple static scenes, we leverage egocentric perception and propose collision-avoiding motion primitives (CAMPs) to enable virtual humans to self-explore and navigate in a dynamic environment. Specifically, CAMPs are trained jointly to produce collision-free motion sequences. At each timestep t, the agent observes the state $\mathbf{s}_t$, performs an action $\mathbf{a}_t$, and receives a reward $r_t = r(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1})$, where $\mathbf{s}_{t+1}$ represents the next state of the environment due to $\mathbf{a}_t$.

Egocentric Sensing As Depth Proxy. We aim to sample actions given by a policy to synthesize realistic human motions. Egocentric perception-driven motion synthesis should arguably use egocentric vision as input. However, depth rendering is costly and RL requires billions of samples to converge [3, 103]. Besides, directly training RL with visual data can be unstable [122]. We thereby use a cheap-to-compute *egocentric sensing* $\mathcal{E}_t$ as a proxy for depth images as illustrated in Fig. 2. N rays are cast evenly from the midpoint of two eyeballs, i.e., the location of the egocentric camera. The field of view $[\theta_{min}, \theta_{max}]$, centered on the 2D projection of the viewing direction $\vec{\mathbf{v}}$, limits the agent's perception to the front area. Rays stop at collisions, with collision detection in 2D. See more details in Supp. Mat.

Agent Representation. The agent is a virtual human represented by an SMPL-X mesh [63]. We further compact the body representation by selecting $M = 67$ body surface markers $\mathbf{x} \in \mathbb{R}^{M\times 3}$ on the mesh following [118].

Environment. Inspired by [108], we aim to learn a library of composable CAMPs. We implement a *finite-horizon* environment based on the generative motion primitive model from [117]. Specifically, a motion primitive is defined as a 0.5-second motion clip containing $T = 20$ frames in the canonical coordinate, and each frame contains a single agent representation. The primitive model $\mathcal{P}$ is based on the C-VAE framework [88], which takes the first $T_s = 2$ frames as the condition, and models a conditional probability of the next $T - T_s$ frames. Compared to [117] trained on the AMASS dataset [48] with many sport motion sequences, we train $\mathcal{P}$ using the SAMP dataset [29], which focuses on daily activities, better suited for HMD use cases. Our *action space* $\mathcal{A}$ is the pretrained 128D latent space of $\mathcal{P}$, and the

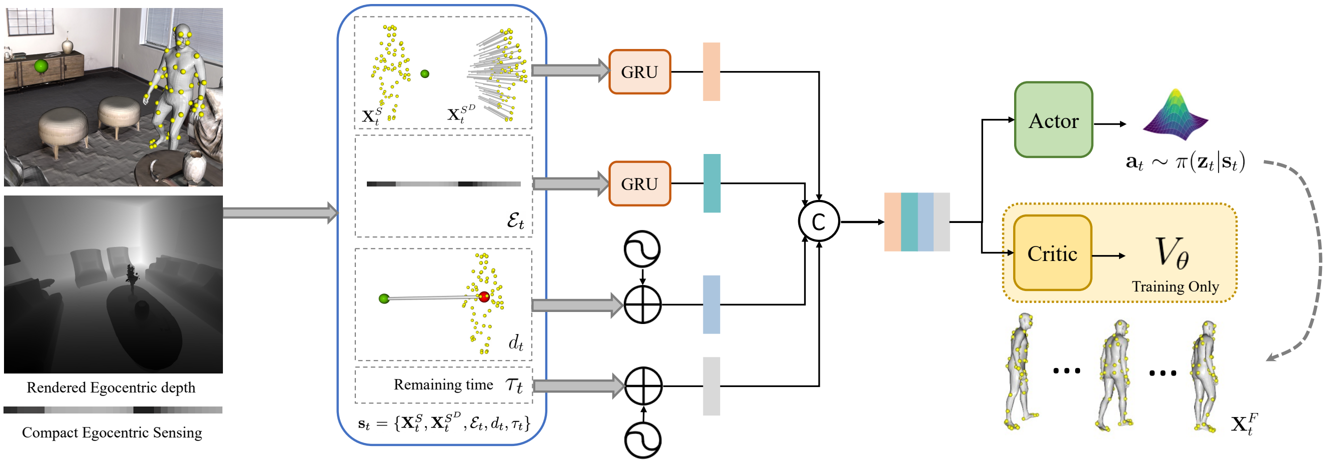

Figure 2. Policy network architecture. We learn a generalizable mapping from motion seed body markers $\mathbf{X}_t^S$, marker directions $\mathbf{X}_t^{S^D}$, egocentric sensing $\mathcal{E}_t$, and distance to the target d_t to CAMPs. The policy learns a stochastic collision avoiding action space to predict future body markers $\mathbf{X}_t^F$. For illustration purposes, we visualize only one frame of $\mathbf{X}_t^S$ and $\mathcal{E}_t$. See Sec. 3.1 and 3.2 for details.

action $\mathbf{a}_t$ can be randomly sampled from $\mathcal{A}$.

With the input of a random action $\mathbf{a}_t \in \mathbb{R}^{128}$ and a motion seed $\mathbf{X}_t^S = [\mathbf{x}_t^0, \mathbf{x}_t^{T_s-1}]$ (history frames), $\mathcal{P}$ predicts future frames $\mathbf{X}_t^F = [\mathbf{x}_t^{T_s}, ..., \mathbf{x}_t^{T-1}]$ of the current motion primitive $\mathbf{X}_t = [\mathbf{X}_t^S, \mathbf{X}_t^F] \in \mathbb{R}^{T\times M\times 3}$, which represents a short sequence of human motion spanning 0.5 s:

$$\mathbf{X}_t^F = \mathcal{P}(\mathbf{X}_t^S, \mathbf{a}_t)$$

State. To preserve Markov property [62], the state is defined as $\mathbf{s}_t = \{\mathbf{X}_t^S, \mathbf{X}_t^{S^D}, \mathcal{E}_t, d_t, \tau_t\}$, in which $\mathbf{X}_t^{S^D} \in \mathbb{R}^{T_s\times M\times 3}$ denotes the normalized direction of each marker seed to the target, $\mathcal{E}_t \in \mathbb{R}^{T_s\times N}$ denotes the egocentric sensing depth proxy, d_t denotes the distance from the pelvis to the target, and τ_t denotes the remaining time. See Fig. 2.

Reward. To synthesize egocentric perception behaviors, we use an "attention reward" to incentivize the virtual human to look in specific directions: $r_{attention} = \cos\langle\vec{\mathbf{v}}, \vec{\mathbf{a}}\rangle$, where $\vec{\mathbf{a}}$ is the attention direction from the head joint to the viewing target. The reward function is defined as:

$$r_t = r_{cont.} + r_{dist} + r_{ori} + r_{attention} + r_{pene} + r_{pose} + r_{succ},$$

where $r_{cont.}$ enforces valid foot contact and minimizes foot skating; r_{dist} encourages reaching the target; r_{ori} aligns the body forward direction with the target; r_{pene} guides collision avoidance; r_{pose} reduces unrealistic human poses; and r_{succ} is a sparse reward when reaching the target.

Episode Termination. To handle collisions beyond [119], we employ multiple *termination* signals to conclude an episode if the generated motion primitive $\mathbf{X}_t$ satisfies any of the following conditions:

- Success: The virtual human reached the target.
- Penetration: The virtual human collides with the obstacle.
- Timeout: The virtual human did not reach the target within the maximum timesteps.

3.2. Training Collision-Avoiding Stochastic Policies

Algorithm. We use Proximal Policy Optimization [81] (PPO) to learn a generalizable mapping from various egocentric sensing and body configurations to CAMPs. Instead of extensive manual data collection for all possible input combinations, we leverage the stochastic nature of the PPO policy. Through exploration and sampling actions, the agent traverses the scene and generates varied egocentric sensing and body configurations, diversifying the training data.

Instead of training each CAMP independently for every single step, we use PPO to train a sequence of CAMPs jointly in multi-step collision avoidance tasks. This approach can benefit choosing a more favorable CAMP which makes the subsequent action easier.

Network. The network architecture is shown in Fig 2. The actor and critic network share a feature extraction trunk to encode the state $\mathbf{s}_t$: the motion seed ($\mathbf{X}_t^S$ and $\mathbf{X}_t^{S^D}$) and the egocentric sensing $\mathcal{E}_t$ are encoded using RNNs; the rest of scalar states are encoded using positional encoding [99]. The actor predicts a stochastic policy $\mathbf{a}_t \sim \pi(\mathbf{z}_t|\mathbf{s}_t)$ conditioned on the current state $\mathbf{s}_t$, where $\mathbf{z}_t \sim \mathcal{N}(\mu, \Sigma)$. μ and Σ are the mean and variance of the learned action space.

Objective Function. The objective function includes the policy surrogate L^{CLIP}, the value function error term L^{VF} to evaluate the value prediction V_θ, and an entropy bonus L^S to encourage exploration:

$$L = L^{CLIP} + c_1 L^{VF} + c_2 L^S$$

where c_1, c_2 are coefficients. See more details in Supp. Mat.

RL Pre-training and Finetuning. Training in crowded scenes, e.g. Replica [91], requires additional steps. Because the action space $\mathcal{A}$ is an unbounded Gaussian, RL exploration while predicting reasonable human poses can

Algorithm 1 Crowd motion synthesis with learned CAMPs

Result: Multi-human locomotion w/ collision avoidance;
Init: crowd size C, marker seed for each human $\mathbf{X}_c^S$;
for $step \leftarrow 1$ **to** max_steps **do** ▷ env. finite horizon
  for $c \leftarrow 1$ **to** C **do** ▷ for each human
    update all locations with $\{bbox(\mathbf{X}_c^S)\}_{c=1:C}$
    compute egocentric sensing $\mathcal{E}_c$;
    execute the action that maximizes the expected return, and produce one CAMP;
  end for
end for

be challenging. We first pretrain the policy with a higher r_{pene} weight without *penetration termination*. After convergence, we finetune it with strict termination constraints using a signed distance field (SDF) for penetration detection. Please refer to Supp. Mat. for the formulation and weighting of each reward and training detail.

3.3. Compositing Learned Motion Primitives

Although CAMPs are trained solely with static scenes, their direct application to dynamic settings is achieved by decomposing jointly trained CAMPs into individual motion primitives and re-compositing them. Our model demonstrates effective generalization by selecting the next best motion primitive from the learned CAMP library to maximize the expected return, provided that the egocentric sensing is updated with the most recent obstacle location at each timestep. Furthermore, our model is directly applicable to tasks involving complex interactions with other virtual humans. To synthesize crowd motion (Alg. 1), each virtual human employs the same policy to navigate and avoid others. To a specific virtual human, others are seen as dynamic obstacles, represented by body bounding boxes for avoidance. Acknowledging the inherent delay in human reactions when avoiding dynamic obstacles [4], agents take a single CAMP sequentially, instead of in parallel, i.e. the first agent generates its first CAMP and waits for others to complete their first CAMP before all agents move on to prepare their second CAMP. To ensure successful collision avoidance, the agent's egocentric sensing is updated before taking a new action. This composition of CAMPs synthesizes emergent multi-human behaviors *without* multi-agent RL algorithms (see Sec. 5.1), enhancing the generalization and scalability.

4. Egocentric Synthetic Data Generation

Synthesizing realistic egocentric perception-driven human motions (as detailed in Sec. 3) forms the foundation of simulating egocentric synthetic data. An overview of our egocentric data generation pipeline *EgoGen*, is shown in Fig. 3.

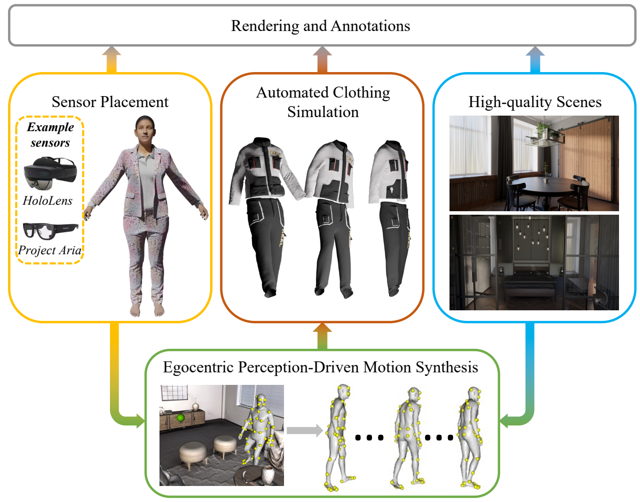

Figure 3. Overview of *EgoGen*. Through generative motion synthesis (Sec. 3), we further enhance egocentric data diversity by randomly sampling diverse body textures (ethnicity, gender) and 3D textured clothing through an automated clothing simulation pipeline (Sec. 4.2). With high-quality scenes and different egocentric cameras, we can render photorealistic egocentric synthetic data with rich and accurate ground truth annotations (Sec. 4.3).

4.1. Embodied Camera Placement

Similar to existing AR devices, we use the head pose to define the egocentric viewing direction $\vec{\mathbf{v}}$. Our development is based on Blender [11]. We use the SMPL-X [63] mesh to position the egocentric camera between the two eyeballs. The camera's viewing direction ($\vec{\mathbf{v}}$) is perpendicular to the plane determined by the two eye bones in the armature. We also support simulating multi-camera rigs as shown in Fig. 1. When the body moves (Sec. 3.3), we can synthesize egocentric videos with continuously updated camera poses.

4.2. Body Texture and Clothing

To enhance *EgoGen*'s synthetic data realism, we dress virtual humans using human textures and 3D clothing assets from BEDLAM [1, 8], including 50 male and 50 female skin albedo textures from seven ethnic groups.

Unlike prior works [8, 109] relying on unscalable commercial software for clothing dynamics simulation, we automate it for diverse synthesized motions and body shapes, minimizing manual effort. Each garment mesh is in a consistent rest pose, i.e., A-Pose (See Fig. 3 middle-left). For each motion sequence, we first repose it to match the body pose in the first frame using linear blend skinning. This involves initializing the clothing geometry by sampling pose and shape blend shapes, along with skinning weights from the nearest multiple SMPL-X vertices in A-Pose. Then we simulate upper and lower garments separately using a state-of-the-art clothing simulation network [24].

Table 1. Evaluation of motion synthesis in scenes with moving obstacles, multiple humans, and path diversity. ↓: lower is better; ↑: higher is better. The best results in each scenario are in boldface. * denotes an improved version for fair comparison. (Sec. 5.1)

Evaluation	Metrics	GAMMA* [117]	DIMOS* [119]	Ours
Mov. obs.	SR (%) ↑	96	83	**100**
	Dist. (m) ↓	0.29	0.55	**0.06**
	Cont. ↑	0.95	0.96	**0.97**
	Pene-S. (%) ↓	9.2	8.4	**3.4**
2 humans	SR (%) ↑	95	88	**100**
	Dist. (m) ↓	0.32	0.41	**0.07**
	Cont. ↑	0.96	**0.98**	0.97
	Pene-H. ↓	27.6	10.7	**0**
4 humans	SR (%) ↑	92	70	**100**
	Dist. (m) ↓	0.41	0.79	**0.07**
	Cont. ↑	0.94	0.95	**0.96**
	Pene-H. ↓	60.4	41.7	**0**
Diversity	SR (%) ↑	96	84	**97**
	Std Dev ↑	0.987	1.05	**1.21**

4.3. Rendering and Annotations

EgoGen supports simulating diverse head-mounted devices with different camera models, such as fisheye and pinhole cameras. Given the camera's intrinsic parameters and relative poses within the camera rig, we can simulate AR devices like Project Aria glasses [50] and HoloLens [51], facilitating synthetic data generation for real-world applications. Camera extrinsic is determined by our generative human motion model. We use Blender [11] to render photorealistic egocentric image sequences with motion blur. We also render out a rich set of ground truth annotations, such as depth maps, surface normals, segmentation masks, world positions, optical flow, etc for egocentric perception tasks.

5. Experiments

We assess the motion quality, generalizability, and diversity of our motion model, highlighting its ability to generalize to unseen complex tasks and comparing it with recent baselines (Sec. 5.1). We evaluate our proposed egocentric sensing as a depth proxy for enhancing agent exploration (Sec. 5.2) and conduct ablation studies (Sec. 5.3).

We further demonstrate the effectiveness of *EgoGen* on three egocentric computer vision tasks in Sec. 5.4, and 5.5. By incorporating synthesized egocentric images, we can enhance the performance of the state-of-the-art algorithms.

5.1. Evaluation of Learned CAMPs

We assess CAMPs' generalizability in dynamic scenes, including scenes with moving obstacles and scenes with multiple individuals. In tests with moving obstacles, the obstacle blocks the person's path by moving between the person and the goal. In multiple human test scenes, lines from their starting and goal locations intersect in the middle, requiring solving human-human penetrations. See detail in Sup. Mat.

In Tab. 1, we compare goal-reaching behaviors with two recent baselines: GAMMA [117] and DIMOS [119]. Baseline methods use navigation meshes and path planning for static scenes, while CAMPs can autonomously avoid dynamic obstacles (Sec. 3.3). For fair comparison in dynamic scenes, we extend the baselines by updating navigation meshes and performing on-the-fly path planning at each time step. The tree-based search as in [117] is disabled for all the methods. **Metrics**: (1) *SR*: Success rate for reaching the goal location within a 0.3m threshold. (2) *Dist.*: Average distance of the final pelvis location to the goal. (3) *Cont.*: The contact metric [117] that measures foot-floor contact and foot skating. (4) *Pene-S.*: Percentage of frames with detected human-scene penetration in moving obstacle scenes. (5) *Pene-H.*: Accurate human-human penetration evaluation metric using COAP [52] in multiple human scenes. Please refer to Supp. Mat. for metric details.

CAMPs outperform the two baselines in dynamic scenarios with moving obstacles and multiple humans, exhibiting lower human-scene and human-human penetrations and a higher goal-reaching success rate. In multiple human scenarios, we observe that in the baselines, dynamically redoing path planning for each human independently can not effectively solve human-human penetration. In contrast, composable CAMPs can generalize well in dynamic settings without using multi-agent RL to synthesize crowd motions.

We assess walking path diversity using the standard deviation of pelvis locations for the same start-target pairs in scenes with a single static box obstacle. As shown in Tab. 1 (Diversity), our approach does not require a pre-computed global path and allows agents to self-explore without being constrained by predefined paths, achieving higher walking path diversity and success rate. This fosters diverse synthetic data generation via more diverse synthesized motion.

Table 2. Evaluation of egocentric sensing. (Sec. 5.2)

Method (sensing range)	SR (%) ↑	Dist. (m) ↓
Local map [119] (0.8 m)	78	0.35
Local map* (7 m)	4	3.04
Egocentric sensing (ours) (7 m)	**95**	**0.12**

5.2. Evaluation of Egocentric Sensing

We assess the exploration ability of our egocentric sensing $\mathcal{E}_t$ in Replica [91] scenes. In Tab. 2, we replace $\mathcal{E}_t$ with a local map [119] in our state $\mathbf{s}_t$, following their encoding method. Relying on local information can trap agents in local optima, e.g., walls beyond their sensing range, resulting in lower SR. Our egocentric sensing acts as a depth proxy, allowing the agent to avoid local optima, explore more effectively than local maps [119] or scandots [3], and achieve higher SR. In addition, our compact representation is more scalable as the sensing range increases, while quadratic local map growth can hinder the policy network's learning.

Table 3. Ablation studies. *Note: in our observation,* $\|VP\|_2 > 15$ *indicates abnormal human poses.* (Sec. 5.3)

	SR (%) ↑	$\|\text{VP}\|_2$ ↓	$\cos(\vec{\mathbf{v}}, \vec{\mathbf{a}})$ ↑
Egocentric depth	8	13.64	0.049
No pretraining	90	28.77	0.918
No attention reward	90	12.26	0.891
Our policy	**92**	**10.57**	**0.940**

5.3. Ablation Studies

We compare our policy with several ablations in Tab. 3:
Egocentric depth: an ablation training an egocentric depth image-based policy without the depth sensing proxy. Egocentric depth images are encoded with a CNN;
No pretraining: an ablation training collision avoidance in crowded scenes with strict penetration termination directly;
No attention reward: an ablation for the viewing direction.

We assess pose naturalness with the maximum pose embedding norm encoded with VPoser [63] and evaluate the attention reward with the cosine similarity between the viewing direction $\vec{\mathbf{v}}$ and the attention direction $\vec{\mathbf{a}}$ (Sec. 3.1).

Directly training RL with egocentric depth images is ineffective due to our high-dimensional action space, emphasizing the value of the compact egocentric sensing representation. Training agents with strict penetration constraints in crowded scenes directly can result in exploring unreasonable action subspaces, given its unbounded Gaussian nature, leading to unrealistic human poses, highlighting the effectiveness of our two-stage RL training scheme. Without the attention reward, the virtual human's capability to attend to a specific direction decreases. All ablation studies are evaluated in Replica. See visuals in Supp. Vid. and Supp. Mat.

5.4. Mapping, Localization, and Tracking for HMD

Mapping and localization. LaMAR [76] is the first mapping and localization benchmark dataset for AR in large-scale scenes. Despite over a year of extensive data collection, the dataset still lacks exhaustive scene coverage, especially in large open spaces. *EgoGen* can let virtual humans explore large-scale scenes, render dense egocentric views, and build a more complete SfM map by extracting image feature points with SuperPoint [15] and matching images with SuperGlue [75]. Despite synthetic images being noisier due to scene quality, SuperGlue [75] matching can filter out noisy feature points and yield reliable matches.

In Tab. 4, we evaluate *EgoGen* by assessing the localization recall at $(1^\circ, 10cm)$ on the validation set in a lobby of ∼120 sqm of the LaMAR CAB location. In addition, we report the number of triangulated 3D points (#P3D) and track length. *EgoGen* improves the 3D reconstruction by yielding more points for a slightly improved track length and also a significantly better localization performance compared to using the real data only. Ng et al. [58] augments mapping images by perturbing real-world camera poses with noise, which may generate unrealistic camera poses (e.g., stuck in a wall or facing the ceiling), limiting egocentric localization effectiveness. Their method also assumes the availability of initial camera poses, which may not always be feasible. In contrast, *EgoGen* augments by virtual humans *randomly* exploring scenes. Our approach holds promise for creating AR mapping and localization datasets for digital twin scenes without manual data collection, providing enhanced privacy preservation, e.g. no need for anonymization. Refer to Supp. Mat. for visualization and implementation details.

Table 4. Mapping and localization evaluation. We augment LaMAR with the same amount of images (248 frames) and report the localization recall at $(1^\circ, 10cm)$ on the validation set. *EgoGen* achieves the highest track length and recall. (Sec. 5.4)

	#P3D ↑	Track length ↑	Recall (%) ↑
LaMAR	1929739	5.1946	66.9
Ng et al. [58]	**1937758**	5.1940	74.9
EgoGen	1936169	**5.2105**	**76.7**

Table 5. Egocentric camera tracking evaluation of models trained with and without synthetic data from *EgoGen*. (Sec. 5.4)

	Pose ↓	Rotation ↓	Transl (mm) ↓
Scratch	1.83	0.74	**1303**
+ *EgoGen* pretrain	**1.67**	**0.62**	1305

Egocentric camera tracking. Egocentric camera tracking for HMD aims to yield device pose trajectories in 3D scenes given egocentric video observations. Recovering camera poses from monocular RGB videos using SLAM [93] is a challenging and ill-posed problem due to scale ambiguity. EgoEgo [43] leverages the knowledge of human motion to address the egocentric HMD tracking problem. Specifically, EgoEgo trains a neural network to infer the translation scaling and rotations from egocentric videos, which improves the HMD tracking performance. However, training this model requires jointly captured data of ground truth HMD trajectories and egocentric videos, which are costly to collect. We address this limitation by using *EgoGen* to synthesize quantities of egocentric videos with accurate camera trajectories to pretrain the model, which proves to improve the tracking performance on real data. We conduct experiments on the GIMO [121] dataset that contains ∼200 short sequences of paired motion-video data in 19 scenes. Using *EgoGen*, we synthesize ∼4k sequences of human movements in their scenes and render corresponding egocentric videos using the same camera intrinsic as GIMO and the embodied camera placement described in Sec. 4.1. We also slightly perturb the camera placement location and orientation to simulate the diversity of how people wear HMDs in real data and avoid overfitting to one specific camera placement. We first pretrain the model with synthetic data generated by *EgoGen*, then finetune it on the real GIMO data.

Tab. 5 shows the egocentric camera tracking performances for models trained with and without synthetic data. Definitions of evaluation metrics can be found in Supp. Mat. The finetuned model benefits from *EgoGen* synthetic data and predicts more accurate camera poses compared to the model trained using real data only.

5.5. Human Mesh Recovery from Egocentric views

Human mesh recovery (HMR) is the key to human behavior understanding from the egocentric view, thus crucial for applications in robotics and AR/VR. Given an egocentric RGB or depth image of a target subject, HMR aims to reconstruct the subject's 3D body pose and shape. However, acquiring and annotating real-world data is expensive, demanding, and time-consuming, with egocentric data being particularly scarce. EgoBody [115] is a recent egocentric dataset capturing two-people interactions, with egocentric depth/RGB frames annotated with SMPL-X body meshes. EgoBody provides ∼180k egocentric RGB frames, and merely ∼23k depth frames due to the low frame rate of the depth sensor, with ∼90k/∼10k in the RGB/depth training set. Such limited data is insufficient to train a learning-based model from scratch. In contrast, with *EgoGen*, large-scale synthetic egocentric data can be generated in a time-efficient way. We leverage *EgoGen* to generate quantities of training frames (300k RGB, 105k depth) of humans moving in EgoBody 3D scenes, rendered from the egocentric view, and annotated by SMPL-X parameters of the target subject. Specifically, RGB images are rendered with lifelike human body textures and 3D clothing, with random lighting.

With the recent HMR regressor, ProHMR [41], we show that pre-training with our synthetic data from *EgoGen* enhances the existing method's capability to generalize on real-world scenarios. Evaluated on the real-world EgoBody test set, we compare two training schemes: (1) trained from scratch on the real-world EgoBody training set ("-scratch"), and (2) pre-trained on synthetic data from *EgoGen* and finetuned on the real-world EgoBody training set ("-ft").

HMR from depth. As no existing methods were proposed for depth-based HMR task, we adapt ProHMR [41] to the depth input by changing the channel number of the first convolution layer. To mimic real-world sensor noise, synthetic noise [27] is added to the rendered depth. G-MPJPE is additionally reported for depth-based HMR as depth images provide global information. As shown in Tab. 6, compared to the model trained only with a limited amount of real-world data (Depth-scratch), errors are significantly reduced for the model pre-trained with our large-scale synthetic data (Depth-ft), in terms of global translation (22.9% lower G-MPJPE), local pose (20.7% lower MPJPE), and body shape (19.5% lower V2V).

HMR from RGB. For training with RGB images, we apply various data augmentation techniques similar to [8].

Table 6. Evaluation of HMR on EgoBody test set. "*-scratch" denotes the model trained from scratch with the Egobody training set, and "*-ft" denotes the model pre-trained with *EgoGen* synthetic data. The units for all metrics are in *mm*. (Sec. 5.5)

	G-MPJPE ↓	MPJPE ↓	PA-MPJPE ↓	V2V ↓
Depth-scratch	117.7	82.2	54.1	100.6
Depth-ft	**90.7**	**65.2**	**47.3**	**81.0**
RGB-scratch	-	90.7	59.9	102.1
RGB-ft	-	**85.3**	**56.2**	**97.2**

Tab. 6 indicates that the RGB-based model pre-trained with large-scale synthetic data ("RGB-ft") also outperforms the model trained only on real-world data ("RGB-scratch"), for both body pose and shape accuracy.

The enhanced performance highlights that *EgoGen*'s synthetic data effectively compensates for the lack of real-world training data, boosting the performance of current methods when test on real-world data. We will release both of our synthetic EgoBody datasets. See dataset statistics, qualitative visualizations, and training details in Supp. Mat.

6. Conclusion and Future Work

We propose a novel egocentric synthetic data generation approach, *EgoGen*, that uses embodied sensors, a parametric body model, and a generative egocentric perception-driven human motion synthesis method to create egocentric training data with accurate and rich ground truth annotations. By integrating deep reinforcement learning and collision-avoiding motion primitives with egocentric depth proxy, *EgoGen* synthesizes robust human motion and emergent multi-agent behaviors. This paves the way to an efficient and scalable data generation solution that may have a profound impact on egocentric perception tasks.

Human-scene interaction in *EgoGen* is currently coarse. We aim to extend the current method to simulate more detailed human motion driven by egocentric perception, such as hand manipulation, sitting, lying, etc, to facilitate more realistic egocentric synthetic data. We use fixed attention goals to model human attention. Predicting human intention through historical egocentric perception and synthesizing viewing directions based on predicted intention holds significant potential. Synthesizing gaze direction for predicting human intent is valuable but presently hampered by data requirements; we will revisit this when resources allow.

We will explore many other egocentric vision tasks with *EgoGen* as this area grows rapidly such as social understanding and forecasting. *EgoGen* could benefit human-robot interaction, e.g., our generative human motion model and lifelike human appearances can be integrated into [67] to close the sim2real gap for robotic agents further.

Acknowledgements. This work is supported by the SNSF project grant 200021 204840 and the SDSC PhD fellowship. We sincerely thank Linfei Pan, Yunke Ao, Qianli Ma, and Maxime Raafat for the valuable discussions.

References

[1] Meshcapade GmbH, Tübingen, Germany. https://meshcapade.com,, 2022. 5

[2] Mohammed Abdel-Wahab, Konrad Wenzel, and Dieter Fritsch. Efficient reconstruction of large unordered image datasets for high accuracy photogrammetric applications. In *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Melbourne, Australia. XXII ISPRS Congress*, 2012. 3

[3] Ananye Agarwal, Ashish Kumar, Jitendra Malik, and Deepak Pathak. Legged locomotion in challenging terrains using egocentric vision. In *Conference on Robot Learning, CoRL 2022, 14-18 December 2022, Auckland, New Zealand*, pages 403–415. PMLR, 2022. 3, 6

[4] M. Pilar Aivar, Eli Brenner, and Jeroen B. J. Smeets. Avoiding moving obstacles. *Experimental Brain Research*, 190 (3):251–264, 2008. 5

[5] Hiroyasu Akada, Jian Wang, Soshi Shimada, Masaki Takahashi, Christian Theobalt, and Vladislav Golyanik. Unrealego: A new dataset for robust egocentric 3d human motion capture. In *European Conference on Computer Vision (ECCV)*, 2022. 2

[6] Apple. ARKit. https://developer.apple.com/arkit/, 2017. 3

[7] Sonia Baltodano, Srinath Sibi, Nikolas Martelaro, Nikhil Gowda, and Wendy Ju. The rrads platform: a real road autonomous driving simulator. In *Proceedings of the 7th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, pages 281–288, 2015. 2

[8] Michael J. Black, Priyanka Patel, Joachim Tesch, and Jinlong Yang. BEDLAM: A synthetic dataset of bodies exhibiting detailed lifelike animated motion. In *Proceedings IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2, 5, 8

[9] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J. Black. Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image. In *European Conference on Computer Vision*, pages 561–578, 2016. 3

[10] Simon Clavet. Motion matching and the road to next-gen animation. In *Proc. of GDC*, 2016. 2

[11] Blender Online Community. *Blender - a 3D modelling and rendering package*. Blender Foundation, Stichting Blender Foundation, Amsterdam, 2018. 5, 6

[12] Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. http://pybullet.org, 2016–2021. 2

[13] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. The epic-kitchens dataset: Collection, challenges and baselines. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 43(11):4125–4141, 2021. 2

[14] Andrew J. Davison. Real-time simultaneous localisation and mapping with a single camera. In *ICCV*, 2003. 3

[15] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description, 2018. 7

[16] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator, 2017. 2

[17] Salehe Erfanian Ebadi, Saurav Dhakad, Sanjay Vishwakarma, Chunpu Wang, You-Cyuan Jhang, Maciek Chociej, Adam Crespi, Alex Thaman, and Sujoy Ganguly. Psp-hdri+: A synthetic dataset generator for pre-training of human-centric computer vision models. In *First Workshop on Pre-training: Perspectives, Pitfalls, and Paths Forward at ICML 2022*, 2022. 2

[18] Matteo Fabbri, Fabio Lanzi, Simone Calderara, Andrea Palazzi, Roberto Vezzani, and Rita Cucchiara. Learning to detect and track visible and occluded body joints in a virtual world. In *European Conference on Computer Vision (ECCV)*, 2018. 2

[19] Alexander Fiannaca, Ilias Apostolopoulous, and Eelke Folmer. Headlock: A wearable navigation aid that helps blind cane users traverse large open spaces. In *Proceedings of the 16th international ACM SIGACCESS conference on Computers & accessibility*, pages 19–26, 2014. 2

[20] Zahra Ghodsi, Siva Kumar Sastry Hari, Iuri Frosio, Timothy Tsai, Alejandro Troccoli, Stephen W. Keckler, Siddharth Garg, and Anima Anandkumar. Generating and characterizing scenarios for safety testing of autonomous vehicles, 2021. 2

[21] Google. ARCore. https://developers.google.com/ar/, 2018. 3

[22] Venu Madhav Govindu. Combining two-view constraints for motion estimation. In *CVPR*, 2001. 3

[23] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, Miguel Martin, Tushar Nagarajan, Ilija Radosavovic, Santhosh Kumar Ramakrishnan, Fiona Ryan, Jayant Sharma, Michael Wray, Mengmeng Xu, Eric Zhongcong Xu, Chen Zhao, Siddhant Bansal, Dhruv Batra, Vincent Cartillier, Sean Crane, Tien Do, Morrie Doulaty, Akshay Erapalli, Christoph Feichtenhofer, Adriano Fragomeni, Qichen Fu, Christian Fuegen, Abrham Gebreselasie, Cristina Gonzalez, James Hillis, Xuhua Huang, Yifei Huang, Wenqi Jia, Weslie Khoo, Jachym Kolar, Satwik Kottur, Anurag Kumar, Federico Landini, Chao Li, Yanghao Li, Zhenqiang Li, Karttikeya Mangalam, Raghava Modhugu, Jonathan Munro, Tullie Murrell, Takumi Nishiyasu, Will Price, Paola Ruiz Puentes, Merey Ramazanova, Leda Sari, Kiran Somasundaram, Audrey Southerland, Yusuke Sugano, Ruijie Tao, Minh Vo, Yuchen Wang, Xindi Wu, Takuma Yagi, Yunyi Zhu, Pablo Arbelaez, David Crandall, Dima Damen, Giovanni Maria Farinella, Bernard Ghanem, Vamsi Krishna Ithapu, C. V. Jawahar, Hanbyul Joo, Kris Kitani, Haizhou Li, Richard Newcombe, Aude Oliva, Hyun Soo Park, James M. Rehg, Yoichi Sato, Jianbo Shi, Mike Zheng Shou, Antonio Torralba, Lorenzo Torresani, Mingfei Yan, and Jitendra Malik. Ego4d: Around the World in 3,000 Hours of Egocentric

Video. In *IEEE/CVF Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2

[24] Artur Grigorev, Michael J. Black, and Otmar Hilliges. HOOD: hierarchical graphs for generalized modelling of clothing dynamics. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 16965–16974. IEEE, 2023. 5

[25] Vladimir Guzov, Aymen Mir, Torsten Sattler, and Gerard Pons-Moll. Human poseitioning system (hps): 3d human pose estimation and self-localization in large scenes from body-mounted sensors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4318–4329, 2021. 3

[26] Norbert Haala, Michael Cramer, Florian Weimer, and Martin Trittler. Performance test on uav-based photogrammetric data collection. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 2012. 3

[27] Ankur Handa, Thomas Whelan, John McDonald, and Andrew J Davison. A benchmark for rgb-d visual odometry, 3d reconstruction and slam. *ICRA*, 2014. 8

[28] Albert Haque, Boya Peng, Zelun Luo, Alexandre Alahi, Serena Yeung, and Li Fei-Fei. Towards viewpoint invariant 3d human pose estimation. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, pages 160–177. Springer, 2016. 3

[29] Mohamed Hassan, Duygu Ceylan, Ruben Villegas, Jun Saito, Jimei Yang, Yi Zhou, and Michael Black. Stochastic scene-aware motion prediction. In *Proceedings of the International Conference on Computer Vision 2021*, 2021. 2, 3

[30] Mohamed Hassan, Yunrong Guo, Tingwu Wang, Michael Black, Sanja Fidler, and Xue Bin Peng. Synthesizing physical character-scene interactions. *arXiv preprint arXiv:2302.00883*, 2023. 2

[31] Daniel Holden, Taku Komura, and Jun Saito. Phase-functioned neural networks for character control. *ACM Trans. Graph.*, 36(4):1–13, 2017. 2

[32] Daniel Holden, Oussama Kanoun, Maksym Perepichka, and Tiberiu Popa. Learned motion matching. *ACM Transactions on Graphics (TOG)*, 39(4):53–1, 2020. 2

[33] Yuhe Jin, Dmytro Mishkin, Anastasiia Mishchuk, Jiri Matas, Pascal Fua, Kwang Moo Yi, and Eduard Trulls. Image Matching across Wide Baselines: From Paper to Practice. *IJCV*, 2021. 3

[34] Neena Kamath. Announcing Azure Spatial Anchors for collaborative, cross-platform mixed reality apps. `https://azure.microsoft.com/en-us/blog/announcing-azure-spatial-anchors-for-collaborative-cross-platform-mixed-reality-apps/`, 2019. 3

[35] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7122–7131, 2018. 3

[36] Alex Kendall, Matthew Grimes, and Roberto Cipolla. PoseNet: A Convolutional Network for Real-Time 6-DoF Camera Relocalization. In *ICCV*, 2015. 3

[37] Georg Klein and David Murray. Parallel tracking and mapping for small ar workspaces. In *IEEE and ACM International Symposium on Mixed and Augmented Reality*, 2007. 3

[38] Muhammed Kocabas, Chun-Hao P. Huang, Otmar Hilliges, and Michael J. Black. PARE: Part attention regressor for 3D human body estimation. In *Proceedings International Conference on Computer Vision (ICCV)*, pages 11127–11137. IEEE, 2021. 3

[39] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3D human pose and shape via model-fitting in the loop. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2252–2261, 2019.

[40] Nikos Kolotouros, Georgios Pavlakos, and Kostas Daniilidis. Convolutional mesh regression for single-image human shape reconstruction. In *CVPR*, 2019.

[41] Nikos Kolotouros, Georgios Pavlakos, Dinesh Jayaraman, and Kostas Daniilidis. Probabilistic modeling for human mesh recovery. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 11605–11614, 2021. 3, 8

[42] Lucas Kovar, Michael Gleicher, and Frédéric H. Pighin. Motion graphs. In *International Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 2008, Los Angeles, California, USA, August 11-15, 2008, Classes*, pages 51:1–51:10. ACM, 2008. 2

[43] Jiaman Li, C. Karen Liu, and Jiajun Wu. Ego-body pose estimation via ego-head pose estimation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 17142–17151. IEEE, 2023. 1, 2, 7

[44] Zhihao Li, Jianzhuang Liu, Zhensong Zhang, Songcen Xu, and Youliang Yan. CLIFF: Carrying location information in full frames into human pose and shape estimation. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part V*, pages 590–606. Springer, 2022. 3

[45] Hung Yu Ling, Fabio Zinno, George Cheng, and Michiel van de Panne. Character controllers using motion vaes. *ACM Trans. Graph.*, 39(4), 2020. 3

[46] Miao Liu, Dexin Yang, Yan Zhang, Zhaopeng Cui, James M Rehg, and Siyu Tang. 4d human body capture from egocentric video via 3d scene grounding. In *2021 International Conference on 3D Vision (3DV)*, pages 930–939. IEEE, 2021. 3

[47] Zhengyi Luo, Ryo Hachiuma, Ye Yuan, Shun Iwase, and Kris M Kitani. Kinematics-guided reinforcement learning for object-aware 3d ego-pose estimation. *arXiv preprint arXiv:2011.04837*, 2020. 3

[48] Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. Amass: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5442–5451, 2019. 3

[49] Angel Martínez-González, Michael Villamizar, Olivier Canévet, and Jean-Marc Odobez. Residual pose: A decoupled approach for depth-based 3d human pose estimation. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 10313–10318. IEEE, 2020. 3

[50] Meta. Project Aria Glasses. https://www.projectaria.com/, 2023. 6

[51] Microsoft. HoloLens 2. https://www.microsoft.com/en-us/hololens, 2019. 6

[52] Marko Mihajlovic, Shunsuke Saito, Aayush Bansal, Michael Zollhoefer, and Siyu Tang. COAP: Compositional articulated occupancy of people. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022. 6

[53] Aymen Mir, Xavier Puig, Angjoo Kanazawa, and Gerard Pons-Moll. Generating continual human motion in diverse 3d scenes. *CoRR*, abs/2304.02061, 2023. 2

[54] Michael Montemerlo, Nicholas Roy, and Sebastian Thrun. Perspectives on standardization in mobile robot programming: the carnegie mellon navigation (CARMEN) toolkit. In *2003 IEEE/RSJ International Conference on Intelligent Robots and Systems, Las Vegas, Nevada, USA, October 27 - November 1, 2003*, pages 2436–2441. IEEE, 2003. 2

[55] Gyeongsik Moon, Ju Yong Chang, and Kyoung Mu Lee. V2v-posenet: Voxel-to-voxel prediction network for accurate 3d hand and human pose estimation from a single depth map. In *Proceedings of the IEEE conference on computer vision and pattern Recognition*, pages 5079–5088, 2018. 3

[56] Etienne Mouragnon, Maxime Lhuillier, Michel Dhome, Fabien Dekeyser, and Patrick Sayd. Real time localization and 3d reconstruction. In *CVPR*, 2006. 3

[57] Evonne Ng, Donglai Xiang, Hanbyul Joo, and Kristen Grauman. You2me: Inferring body pose in egocentric video via first and second person interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9890–9900, 2020. 3

[58] Tony Ng, Adrian Lopez-Rodriguez, Vassileios Balntas, and Krystian Mikolajczyk. Reassessing the limitations of cnn methods for camera pose regression, 2021. 7

[59] Niantic. Niantic Expands Developer Platform and AR Tools with Niantic Lightship. https://nianticlabs.com/news/lightship/, 2021. 3

[60] David Nistér, Oleg Naroditsky, and James Bergen. Visual odometry. In *CVPR*, 2004. 3

[61] Xiaqing Pan, Nicholas Charron, Yongqian Yang, Scott Peters, Thomas Whelan, Chen Kong, Omkar M. Parkhi, Richard A. Newcombe, and Carl Yuheng Ren. Aria digital twin: A new benchmark dataset for egocentric 3d machine perception. *CoRR*, abs/2306.06362, 2023. 2

[62] Fabio Pardo, Arash Tavakoli, Vitaly Levdik, and Petar Kormushev. Time limits in reinforcement learning. In *Proceedings of the 35th International Conference on Machine Learning*, pages 4045–4054, Stockholmsmässan, Stockholm Sweden, 2018. PMLR. 4

[63] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3D hands, face, and body from a single image. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 10975–10985, 2019. 3, 5, 7

[64] Xue Bin Peng, Ze Ma, Pieter Abbeel, Sergey Levine, and Angjoo Kanazawa. Amp: Adversarial motion priors for stylized physics-based character control. *ACM Transactions on Graphics (TOG)*, 40(4):1–20, 2021. 2

[65] Mathis Petrovich, Michael J. Black, and Gül Varol. Action-conditioned 3d human motion synthesis with transformer vae, 2021. 3

[66] Marc Pollefeys, Luc Van Gool, Maarten Vergauwen, Frank Verbiest, Kurt Cornelis, Jan Tops, and Reinhard Koch. Visual modeling with a hand-held camera. *IJCV*, 2004. 3

[67] Xavier Puig, Eric Undersander, Andrew Szot, Mikael Dallaire Cote, Tsung-Yen Yang, Ruslan Partsey, Ruta Desai, Alexander William Clegg, Michal Hlavac, So Yeon Min, Vladimir Vondrus, Théophile Gervet, Vincent-Pierre Berges, John M. Turner, Oleksandr Maksymets, Zsolt Kira, Mrinal Kalakrishnan, Jitendra Malik, Devendra Singh Chaplot, Unnat Jain, Dhruv Batra, Akshara Rai, and Roozbeh Mottaghi. Habitat 3.0: A co-habitat for humans, avatars and robots. *CoRR*, abs/2310.13724, 2023. 2, 8

[68] Albert Pumarola, Jordi Sanchez, Gary Choi, Alberto Sanfeliu, and Francesc Moreno-Noguer. 3DPeople: Modeling the Geometry of Dressed Humans. In *International Conference in Computer Vision (ICCV)*, 2019. 2

[69] Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, Andrew Y Ng, et al. Ros: an open-source robot operating system. In *ICRA workshop on open source software*, page 5. Kobe, Japan, 2009. 2

[70] Umer Rafi, Juergen Gall, and Bastian Leibe. A semantic occlusion model for human pose estimation from a single depth image. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 67–74, 2015. 3

[71] Tilman Reinhardt. Google Visual Positioning Service. https://ai.googleblog.com/2019/02/using-global-localization-to-improve.html, 2019. 3

[72] Davis Rempe, Tolga Birdal, Aaron Hertzmann, Jimei Yang, Srinath Sridhar, and Leonidas J. Guibas. Humor: 3d human motion model for robust pose estimation. In *International Conference on Computer Vision (ICCV)*, 2021. 3

[73] Davis Rempe, Zhengyi Luo, Xue Bin Peng, Ye Yuan, Kris Kitani, Karsten Kreis, Sanja Fidler, and Or Litany. Trace and pace: Controllable pedestrian animation via guided trajectory diffusion, 2023. 2, 3

[74] Guodong Rong, Byung Hyun Shin, Hadi Tabatabaee, Qiang Lu, Steve Lemke, Mārtiņš Možeiko, Eric Boise, Geehoon Uhm, Mark Gerow, Shalin Mehta, et al. Lgsvl simulator: A high fidelity simulator for autonomous driving. In *2020 IEEE 23rd International conference on intelligent transportation systems (ITSC)*, pages 1–6. IEEE, 2020. 2

[75] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superglue: Learning feature matching with graph neural networks, 2020. 7

[76] Paul-Edouard Sarlin, Mihai Dusmanu, Johannes L. Schönberger, Pablo Speciale, Lukas Gruber, Viktor Larsson, Ondrej Miksik, and Marc Pollefeys. LaMAR: Benchmarking Localization and Mapping for Augmented Reality. In *ECCV*, 2022. 3, 7

[77] Torsten Sattler, Will Maddern, Carl Toft, Akihiko Torii, Lars Hammarstrand, Erik Stenborg, Daniel Safari, Masatoshi Okutomi, Marc Pollefeys, Josef Sivic, Fredrik Kahl, and Tomas Pajdla. Benchmarking 6DoF outdoor visual localization in changing conditions. In *CVPR*, 2018. 3

[78] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, Devi Parikh, and Dhruv Batra. Habitat: A platform for embodied ai research, 2019. 2

[79] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *CVPR*, 2016. 3

[80] Johannes L. Schönberger, Hans Hardmeier, Torsten Sattler, and Marc Pollefeys. Comparative evaluation of handcrafted and learned local features. In *CVPR*, 2017. 3

[81] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017. 4

[82] Akash Sengupta, Ignas Budvytis, and Roberto Cipolla. Synthetic training for accurate 3d human pose and shape estimation in the wild. In *British Machine Vision Conference (BMVC)*, 2020. 2

[83] Takaaki Shiratori, Hyun Soo Park, Leonid Sigal, Yaser Sheikh, and Jessica K. Hodgins. Motion capture from body-mounted cameras. *ACM Trans. Graph.*, 30(4):31, 2011. 3

[84] Jamie Shotton, Ross Girshick, Andrew Fitzgibbon, Toby Sharp, Mat Cook, Mark Finocchio, Richard Moore, Pushmeet Kohli, Antonio Criminisi, Alex Kipman, et al. Efficient human pose estimation from single depth images. *IEEE transactions on pattern analysis and machine intelligence*, 35(12):2821–2840, 2012. 3

[85] Jamie Shotton, Ben Glocker, Christopher Zach, Shahram Izadi, Antonio Criminisi, and Andrew Fitzgibbon. Scene Coordinate Regression Forests for Camera Relocalization in RGB-D Images. In *CVPR*, 2013. 3

[86] Jamie Shotton, Toby Sharp, Alex Kipman, Andrew Fitzgibbon, Mark Finocchio, Andrew Blake, Mat Cook, and Richard Moore. Real-time human pose recognition in parts from single depth images. *Commun. ACM*, 56(1):116–124, 2013. 2

[87] Noah Snavely, Steven M. Seitz, and Richard Szeliski. Photo tourism: exploring photo collections in 3d. *ACM Trans. Graph.*, 25(3):835–846, 2006. 3

[88] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2015. 3

[89] Sebastian Starke, He Zhang, Taku Komura, and Jun Saito. Neural state machine for character-scene interactions. *ACM Trans. Graph.*, 38(6):209–1, 2019. 2

[90] Sebastian Starke, Ian Mason, and Taku Komura. Deepphase: Periodic autoencoders for learning motion phase manifolds. *ACM Transactions on Graphics (TOG)*, 41(4):1–13, 2022. 2

[91] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J. Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, Anton Clarkson, Mingfei Yan, Brian Budge, Yajie Yan, Xiaqing Pan, June Yon, Yuyang Zou, Kimberly Leon, Nigel Carter, Jesus Briales, Tyler Gillingham, Elias Mueggler, Luis Pesqueira, Manolis Savva, Dhruv Batra, Hauke M. Strasdat, Renzo De Nardi, Michael Goesele, Steven Lovegrove, and Richard Newcombe. The Replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019. 4, 6

[92] Andrew Szot, Alex Clegg, Eric Undersander, Erik Wijmans, Yili Zhao, John Turner, Noah Maestre, Mustafa Mukadam, Devendra Chaplot, Oleksandr Maksymets, Aaron Gokaslan, Vladimir Vondrus, Sameer Dharur, Franziska Meier, Wojciech Galuba, Angel Chang, Zsolt Kira, Vladlen Koltun, Jitendra Malik, Manolis Savva, and Dhruv Batra. Habitat 2.0: Training home assistants to rearrange their habitat. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 2

[93] Zachary Teed and Jia Deng. DROID-SLAM: Deep Visual SLAM for Monocular, Stereo, and RGB-D Cameras. *Advances in neural information processing systems*, 2021. 7

[94] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit H Bermano. Human motion diffusion model. *arXiv preprint arXiv:2209.14916*, 2022. 3

[95] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033, 2012. 2

[96] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012. 2

[97] Denis Tome, Patrick Peluse, Lourdes Agapito, and Hernan Badino. xr-egopose: Egocentric 3d human pose from an hmd camera. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7728–7738, 2019. 2, 3

[98] Denis Tome, Thiemo Alldieck, Patrick Peluse, Gerard Pons-Moll, Lourdes Agapito, Hernan Badino, and Fernando De la Torre. Selfpose: 3d egocentric pose estimation from a headset mounted camera. *arXiv preprint arXiv:2011.01519*, 2020. 3

[99] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. 4

[100] Jian Wang, Diogo Luvizon, Weipeng Xu, Lingjie Liu, Kripasindhu Sarkar, and Christian Theobalt. Scene-aware egocentric 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13031–13040, 2023. 3

[101] Keze Wang, Shengfu Zhai, Hui Cheng, Xiaodan Liang, and Liang Lin. Human pose estimation from depth images via

inference embedded multi-task learning. In *Proceedings of the 24th ACM international conference on Multimedia*, pages 1227–1236, 2016. 3

[102] Xin Wang, Taein Kwon, Mahdi Rad, Bowen Pan, Ishani Chakraborty, Sean Andrist, Dan Bohus, Ashley Feniello, Bugra Tekin, Felipe Vieira Frujeri, Neel Joshi, and Marc Pollefeys. Holoassist: an egocentric human interaction dataset for interactive ai assistants in the real world. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 20270–20281, 2023. 2

[103] Erik Wijmans, Abhishek Kadian, Ari Morcos, Stefan Lee, Irfan Essa, Devi Parikh, Manolis Savva, and Dhruv Batra. DD-PPO: learning near-perfect pointgoal navigators from 2.5 billion frames. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. 3

[104] Kyle Wilson and Noah Snavely. Robust global translations with 1dsfm. In *ECCV*, 2014. 3

[105] Jungdam Won, Deepak Gopinath, and Jessica Hodgins. Physics-based character controllers using conditional vaes. *ACM Trans. Graph.*, 41(4), 2022. 2

[106] Erroll Wood, Tadas Baltrušaitis, Charlie Hewitt, Sebastian Dziadzio, Thomas J Cashman, and Jamie Shotton. Fake it till you make it: face analysis in the wild using synthetic data alone. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3681–3691, 2021. 2

[107] Fu Xiong, Boshen Zhang, Yang Xiao, Zhiguo Cao, Taidong Yu, Joey Tianyi Zhou, and Junsong Yuan. A2j: Anchor-to-joint regression network for 3d articulated pose estimation from a single depth image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 793–802, 2019. 3

[108] Mengyuan Yan, Gen Li, Yilin Zhu, and Jeannette Bohg. Learning topological motion primitives for knot planning. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2020, Las Vegas, NV, USA, October 24, 2020 - January 24, 2021*, pages 9457–9464. IEEE, 2020. 3

[109] Zhitao Yang, Zhongang Cai, Haiyi Mei, Shuai Liu, Zhaoxi Chen, Weiye Xiao, Yukun Wei, Zhongfei Qing, Chen Wei, Bo Dai, Wayne Wu, Chen Qian, Dahua Lin, Ziwei Liu, and Lei Yang. Synbody: Synthetic dataset with layered human models for 3d human perception and modeling, 2023. 5

[110] Heyuan Yao, Zhenhua Song, Baoquan Chen, and Libin Liu. Controlvae: Model-based learning of generative controllers for physics-based characters. *ACM Trans. Graph.*, 41(6), 2022. 3

[111] Mao Ye, Xianwang Wang, Ruigang Yang, Liu Ren, and Marc Pollefeys. Accurate 3d pose estimation from a single depth image. In *2011 International Conference on Computer Vision*, pages 731–738. IEEE, 2011. 3

[112] Vickie Ye, Georgios Pavlakos, Jitendra Malik, and Angjoo Kanazawa. Decoupling human and camera motion from videos in the wild. *arXiv preprint arXiv:2302.12827*, 2023. 3

[113] Ye Yuan and Kris Kitani. Ego-pose estimation and forecasting as real-time pd control. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10082–10092, 2019. 3

[114] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiondiffuse: Text-driven human motion generation with diffusion model. *arXiv preprint arXiv:2208.15001*, 2022. 3

[115] Siwei Zhang, Qianli Ma, Yan Zhang, Zhiyin Qian, Taein Kwon, Marc Pollefeys, Federica Bogo, and Siyu Tang. Egobody: Human body shape and motion of interacting people from head-mounted devices. In *European conference on computer vision (ECCV)*, 2022. 1, 2, 8

[116] Siwei Zhang, Qianli Ma, Yan Zhang, Sadegh Aliakbarian, Darren Cosker, and Siyu Tang. Probabilistic human mesh recovery in 3d scenes from egocentric views. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. 3

[117] Yan Zhang and Siyu Tang. The wanderings of odysseus in 3d scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20481–20491, 2022. 3, 6

[118] Yan Zhang, Michael J Black, and Siyu Tang. We are more than our joints: Predicting how 3d bodies move. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3372–3382, 2021. 3

[119] Kaifeng Zhao, Yan Zhang, Shaofei Wang, Thabo Beeler, , and Siyu Tang. Synthesizing diverse human motions in 3d indoor scenes. In *International conference on computer vision (ICCV)*, 2023. 2, 3, 4, 6

[120] Yuhang Zhao, Elizabeth Kupferstein, Brenda Veronica Castro, Steven Feiner, and Shiri Azenkot. Designing ar visualizations to facilitate stair navigation for people with low vision. In *Proceedings of the 32nd annual ACM symposium on user interface software and technology*, pages 387–402, 2019. 2

[121] Yang Zheng, Yanchao Yang, Kaichun Mo, Jiaman Li, Tao Yu, Yebin Liu, C Karen Liu, and Leonidas J Guibas. Gimo: Gaze-informed human motion prediction in context. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XIII*, pages 676–694. Springer, 2022. 7

[122] Ziwen Zhuang, Zipeng Fu, Jianren Wang, Christopher Atkeson, Sören Schwertfeger, Chelsea Finn, and Hang Zhao. Robot parkour learning. In *Conference on Robot Learning (CoRL)*, 2023. 2, 3